DEFENCE DÉFENSE

Automatic Target Recognition Using HNeT

An investigation of Holographic/Quantum Neural Technology

R.A. English

Defence R&D Canada
TECHNICAL MEMORANDUM
DREO TM 2001-080
December 2001

National Defence Défense nationale

Canada

Automatic Target Recognition Using HNeT

An investigation of Holographic/Quantum Neural Technology

R.A. English
Defence Research Establishment Ottawa

Defence Research Establishment Ottawa

Technical Memorandum

DREO TM 2001-080

December 2001

Abstract

With the release of the Moving and Stationary Target Acquisition and Recognition (MSTAR) public data set, high quality Synthetic Aperture Radar (SAR) imagery of military ground vehicles has been made accessible to the entire research community. Furthermore, standard methods for evaluating classifier results on this data set have been created and released. Using these tools, we reconsider a previously contracted application of AND Corporation's Holographic/Quantum Neural Technology (HNeT) classifier, performing brief analyses of the way HNeT selects features for Automatic Target Recognition (ATR) purposes, the methodology used in the contract and their results, as well as obtaining new results that comply with the MSTAR standard evaluation criteria. These results provide measures of performance for the HNeT classifier using Confusion Matrices and Receiver Operating Characteristic (ROC) curves, that are used to compare with the open literature performance of the MSTAR baseline and two other classifiers, from which we conclude that HNeT outperforms the other three and provides improved ATR.

Résumé

Depuis la publication des données publiques du programme MSTAR (acquisition et reconnaissance de cibles mobiles et fixes), les milieux de la recherche ont accès à des images RAS (radar à ouverture synthétique) de grande qualité de véhicules militaires terrestres. En outre, des méthodes d'évaluation uniformes des résultats des classificateurs relativement à cet ensemble de données ont été définies et publiées. À l'aide de ces outils, nous réexaminons l'application du classificateur HNeT (technologie neuronale holographique/quantique) de la corporation AND lors d'un contrat antérieur, exécutant de brèves analyses de la sélection de caractéristiques par HNeT pour la reconnaissance automatique de cible (RAC) ansi que de la méthodologie utilisée et des résultats obtenus lors du contrat. De nouveaux résultats répondant aux critères d'évaluation uniformes de MSTAR sont également obtenus. Ceux-ci fournissent, grâce à l'utilisation de matrices de confusion et de courbes de caractéristique de fonctionnement de récepteur (CFR), une évaluation de la performance du classificateur HNeT et permettent une comparaison avec la performance du classificateur MSTAR de base documentée en littérature ouverte ainsi que de deux autres classificateurs. Nous concluons que HNeT est supérieur aux trois autres classificateurs et offre une RAC améliorée.

This page intentionally left blank.

Executive summary

Automated Target Recognition (ATR) remains a highly desirable, yet largely unfulfilled, goal for defence research. Recently, AND Corporation has commercially released its novel classifier HNeT, upon which an ATR system can be built, and has claimed unprecedented success in HNeT's applications. The author's analysis of the operation of HNeT is presented, explaining how HNeT is able to achieve such a strong performance. In addition to test results generated by AND under DREO contract, the author presents his own application of HNeT that meets the standard evaluation criteria set by DARPA/Wright Laboratory's MSTAR program, applied to SAR images of military vehicles, and shows better results from HNeT than those of the baseline classifier published by MSTAR.

Performance measures of HNeT are given by Confusion Matrices and Receiver Operating Characteristic (ROC) curves tailored for the publicly released MSTAR data collection and compared to classifiers from the open literature, including the MSTAR baseline results. These evaluations show a significant increase in performance by using the HNeT classifier over the others. Specifically, the ROC curves for HNeT reduce the region of error under the MSTAR criteria by about half, as evidenced by an improvement, from 89% correct classification of declared targets for the MSTAR baseline to 95% for HNeT, from which we conclude that HNeT is a strong candidate to provide the basis for an operational ATR capability.

R.A. English (2001). Automatic Target Recognition Using HNeT. DREO TM 2001-080. Defence Research Establishment Ottawa.

Sommaire

La reconnaissance automatique de cible demeure un objectif hautement désirable pour la recherche de défense, mais reste en grande partie à réaliser. Récemment, la société AND a mis en marché son classificateur novateur HNeT, sur lequel un système de reconnaissance automatique de cible (ATR) peut être construit. AND a fait état d'un succès sans précédent en ce qui concerne les applications de HNeT. L'auteur présente une analyse du fonctionnement de HNeT et explique comment ce classificateur est capable d'une telle performance. Outre les résultats d'essai obtenus par AND à contrat pour le CRDO, l'auteur présente sa propre application de HNeT, laquelle répond aux critères d'évaluation uniformes du programme MSTAR du laboratoire Wright de DARPA, appliqués aux images RAS de véhicules militaires, et obtient de HNeT des résultats supérieurs à ceux du classificateur de base publiés par MSTAR.

Des mesures de la performance de HNeT sont obtenues par l'entremise de matrices de confusion et de courbes de caractéristique de fonctionnement de récepteur (CFR), convenablement modifié pour la collection de donnés MSTAR disponible publiquement, et comparées avec des classificateurs disponibles en littérature ouverte incluant les résultats MSTAR de base. Ces évaluations démontrent qu'une importante amélioration de la performance résulte de l'utilisation de classificateur HNeT au lieu des autres classificateurs. Particulièrement, les courbes CFR de HNeT réduisent approximativement de moitié la région d'erreur du critère MSTAR, comme démontré par le passage d'un taux de classification correcte de cibles avouées de 89% pour le critère MSTAR de base à un taux de 95% pour HNeT, ce qui porte à conclure que HNeT est un candidat sérieux pour l'implémentation d'une capabilité RAC opérationnelle.

R.A. English (2001). [Reconnaissance automatique de cible au moyen de HNeT: étude de la technologie neuronale holographique/quantique]. DREO TM 2001-080. Centre de recherches pour la défense Ottawa.

Table of contents

List of figures

List of tables

This page intentionally left blank.

1. Introduction

Research into *Automatic Target Recognition* (ATR) has received a dual boost with recent access to the *Moving and Stationary Target Acquisition and Recognition* (MSTAR) data set of Spotlight *Synthetic Aperture Radar* (SAR) vehicle images which the U.S. *Defense Advanced Research Projects Agency* (DARPA) has made public and with the release of AND Corporation's *Holographic/Quantum Neural Technology* (HNeT).

The public MSTAR data set provides approximately 20,000 SAR images of 10 vehicle types from the former Soviet Union. The vehicles have been imaged over the full 360°azimuthal circle at 1° increments while at several depression angles: 15, 17, 30 and 45 degrees. Three collections have occurred — Fall '95 (Huntsville), Fall '96 (Eglin AFB) and Spring '97 (Eglin AFB) — each providing up to three scenes of different vehicle placements. Two of the vehicle types, the BMP-2 and T-72, have multiple variants/versions in the same scene. Thus, a wide variety of *Extended Operating Conditions* (EOCs) exist upon which to investigate the performance of any given ATR method.

Furthermore, standard SAR ATR evaluation experiments for use with the MSTAR public data have been published [1] by the *Wright Laboratory* (WL) at Wright-Patterson AFB, allowing comparison with their baseline results. This heralds the beginning of efforts to address a serious deficiency in ATR research, namely the lack of a self-consistent theory for ATR or the broader technologies of *pattern recognition* or *machine vision.* Section 2 will examine the current state of the art.

The HNeT classification software has been developed commercially by AND Corporation as a replacement for Neural Network technology. Exploiting the mathematics of coherence, used successfully in holography, quantum mechanics and SAR imagery processing, HNeT is able to efficiently learn, process and store information. Running on a Pentium III 450 MHz, HNeT can process several hundred 64x64 pixel images per second during both the training and classifying processing. Rather than limiting the classifier to using pre-selected features, HNeT manipulates a large potential-feature space using coherence to reinforce features that are common-valued across the in-class training data while random-valued features average to zero. These resulting invariant features form the test criteria used to classify new data. Section 3 describes this methodology in more detail.

In mid-1999, AND Corporation was contracted (W7714-8-0200) to create an interface for preprocessing SAR images so as to allow ATR using their HNeT software. The resulting ANDSAR package was then used to import MSTAR images, and several experiments run to identify optimal HNeT and ANDSAR settings for SAR ATR. In particular, logarithmic scaling of the pixel intensities and rotation to common azimuthal orientation are recommended. Section 4 describes and interprets the results generated by the contracted work.

To verify that the HNeT was indeed performing according to the claims of the company, these experiments were repeated with fully sanitized headers, as well as complete isolation of the training and testing sets. One parameter was, in fact, set to minimize error according to the results of the test set. Fortunately, when this was switched to optimize using the training set results, no noticeable change in performance was observed.

The publication of the MSTAR standard evaluation has allowed direct comparison with the baseline ATR results, given as *Receiver Operating Characteristic* (ROC) curves and Confusion Matrices. In the DARPA/WL three vehicle type experiment using multiple BMP-2, BTR-70 and T-72 vehicles, all BMP-2 and T-72 images are available in the public release. The training set consists of images of one vehicle from each type at 17° depression angle and over the full azimuthal set. The test set images include all vehicles of the three types at 15° depression angle. Confusion matrices are evaluated with the Percent Declaration (P_d) set to 0.9, which forces 10% of the target images to be declared as not one of the target types (i.e., classified as Other). Following the process used by AND in their experiments, HNeT has been evaluated based on the DARPA/ML standards, with comparative results shown in section 5.

2. The ATR Capability

The implementation an ATR system requires three fundamental technologies to perform harmoniously: one or more sensors to collect a signal, the processing/calibration to place the signal into a standard frame of reference, and pattern recognition to locate, extract and identify pieces of the signal.

For this study, the sensor is a single imaging radar operating in X-band with a spotlight mode. The signal processing involves standard SAR techniques to produce 1 ft. resolution imagery. This leaves the pattern recognition technology to be the focus of attention for the current investigation.

There are two fundamental types of recognition problems: open and closed. So far, much of the success in pattern recognition has been achieved in the latter, where there are a fixed finite number of known classes to which everything belongs, with noise being the primary source of ambiguity within the data classes. The open problem is far more difficult, since it requires the creation of an "Other" class to describe everything not in one of the pre-defined classes. By its very nature, the "Other" class is not and cannot be completely defined. Automatic Target Recognition (ATR) presents an open problem, where even the simplest case of a single target class cannot provide a clear delineation between in-class images and everything else.

2.1 Development of ATR Theory

It is very important to understand that classical theories attempt to describe pattern recognition within the structures provided by statistical and probability theories [2, 3]. Most of the successes of these theories are limited to closed problems and do not readily apply to open ones. In particular, it is becoming recognized that the lack of a tractable theory of pattern recognition is a major impediment to the development of many operational recognition systems including for ATR [4, 5, 6, 7]. A theory of ATR would provide a mathematical foundation with which the performance of a given ATR could be rigorously extrapolated to Operating Conditions beyond the scope of a given experiment. Until such a theory is developed, all evidence of success in ATR is entirely anecdotal: the only way to know how well an ATR system will work under new conditions is to try it.

The lack of a mathematical foundation means that there is no ability for prediction. Any promises of operational performance must be treated with a healthy dose of scepticism, because the scientific tools needed to validate those promises do not yet exist. In terms of providing an operational ATR, this means that an autonomous system cannot be employed, because we cannot predict how well it will work. However, this deficiency does not prevent the use of ATR algorithms and techniques for *assisted* target recognition, where the *image analyst* (IA) is able to learn the limitations of the ATR tools through trial and error.

As such, Target Recognition continues to be operationally an Art rather than a Science, relying on the qualitative experiences of the human analysts to determine what and where the limitations of their ATR tools are. In this capacity, namely Assisted Target Recognition as opposed to Automatic Target Recognition, research on ATR systems can continue to improve and provide better tools for Signal and Image Analysts to better perform their arts.

Although individual authors identified the lack of a suitable theory for pattern recognition as far back as 1986, it has been only since the mid-1990s that the research community has begun to collectively recognize and address the inadequacy [8]. Despite this, many researchers choose to ignore the problem of evaluation and so overstate their successes, and will continue to do so until a tractable theory exists. Research into development of a theory specific to SAR ATR is currently being spear-headed under DARPA's MSTAR Performance Estimation Theory (PET) program [9], which was begun in 1998.

2.2 Feature Extraction

A prerequisite to classifying a pattern is feature extraction, where aspects of the pattern are chosen to compare against the definitions of the classes. Ideally, the feature set will contain the fewest features needed to maximize the separation between classes. While this is within the realm of possibilities for closed problems, the lack of a complete

definition of the "Other" class in open problems means that the amount of separation for any feature set is uncertain and will be less than optimal. Which features will separate classes, even which ones are independent, is a function of the problem being considered and there are several approaches to extracting features for recognition purposes.

In the capacity of assisting IAs to perform recognition, it is often useful to select features that match those that the IA will be using to define various classes. In this way, it is straightforward for the human operator to verify the results. Size, shape, relative pixel magnitudes, frequency and velocity are all direct human interpretable features. They are, however, also the features most prominently selected for alteration by countermeasures attempting to obscure a target's identity. The signal may also be transformed to yield characteristics that operators can be trained to identify within the transformation space, and recognition algorithms can operate on features within this transform space. For example, Fourier/Wavelet coefficients [10], Geometric/Zernike moments [11], Object/Image invariants [12] and polarimetric elemental structures [13] can yield characteristic features for pattern recognition.

None of the above groups of processes is likely to provide optimal class separation, however, which usually means a reduced classifier performance should be expected. Several approaches have been adopted and developed to address the problem from the point of view of maximizing the separation and thereby optimizing the classifier performance.

Principal Component Analysis is a classical statistical method used to generate a hierarchy of orthogonal components (features) exhibiting the maximal variation within the current data space [14]. In contrast, the more recent Independent Component Analysis (ICA) seeks to generate components that are independent but with the same goal of maximizing variation of data along these components [15, 16]. HNeT has a similar goal, but seeks to achieve it by determining coherent feature invariants. A large space of potential features is calculated for the training set and combined coherently. Features common to a class will be reinforced, while those that take on random values will average to zero.

2.3 Classifier Training

Once the feature extraction for a given recognition problem has been decided, there are three main learning methods by which the selected features can be presented to the classifier for training.

Model-based training offers the ability to train a classifier without active collection of sensor data for every class to be recognized [17]. Development of appropriate models may, however, be just as or even more labour intensive as data collection, but can be accomplished without the knowledge or irritation of the potential target. By not using sensor-based training, there is a significant potential for introducing model artifacts or failure to include important features that are beyond human recognition.

Library-based training offers a collection of real sensor-based data against which an unknown can be compared. Since the data is unadulterated, all features will be available, human recognizable or not. Furthermore, the best match can be easily displayed to an IA for verification. The library method of training is resource intensive and slow, however, ideally requiring the data to include all conditions that will be encountered in the operational environment.

The template method of training tries to strike a balance between the two, by intelligently combining groups of a data collection to form a template or model representation. Since this is done using real data, templates are less susceptible to artifacts, will contain features beyond human recognition and yet allow interpolation of the data.

2.4 Classification Methods

The most important aspect to recognition is, of course, the classifier or predictor. Strictly speaking, a classifier provides a single best guess as to the target class, while a predictor provides a probability or weighting of the target belonging to each class available. Generically, the term classifier is used to encompass both types of processes, since selecting the highest result only from a predictor will yield a classifier.

Classifiers can also be distinguished by the way they handle a change in the classes defined. Adding a new class may be modular; the classifier need only be trained on the new class in order to operate. Other classifiers require complete retraining, which may be time and resource intensive. By doing so, however, the latter group is potentially able to improve performance on the previous classes by using information from the new one. The addition of a new class to a modular classifier can only ever decrease the performance on the pre-existing classes, even though the overall performance may go up.

The simplest classifier to implement is the use of If-Then Rules, which is the natural decision method for most programming languages. This algorithm is the direct implementation of the methods used by human analysts. Unlike IAs, the rules are applied sequentially whereas a human can process the rules as a collection, i.e. in parallel. This means the Rules-based classifier can provide improved speed and consistency, but can never outperform an IA.

In order to analyze information in parallel, models of the human thought process have been developed: Neural Networks (NN). To be implemented as a classifier, Neural Nets need to be trained to a stable, converged state [18, 19] that is supposed to be attained by using a method such as classical back-propagation, essentially using a feedback loop to iterate to the desired state. Specific techniques to address this problem, such as Adaptive Resonance Theory (ART) have been introduced. However, there is no guarantee that a given system, especially for an open problem, even has a stable, converged state [20]. If not, the NN can still potentially outperform human recognition

when training brings it near an optimized, though unstable, state. Support Vector Machines (SVM) apply statistical analysis to the features for a pair of classes in the system [21], maximizing the separation between the classes in a non-linear fashion. Since the SVM operate as binary classifiers, they are better suited to systems with a small number of classes, but this is merely an issue of resources.

Instead of using pre-defined rules, Decision Trees [14] generate their own, which may include what would be considered collections for If-Then Rule implementation. As such, Decision Trees are capable of mimicking parallel classification, even though rules are applied sequentially.

HNeT performs classification by feature invariant matching, which involves determining and assigning a relative weighting to features that recur in the in-class training data and not in the outclass data. The features of incoming data are then matched against the various class invariants.

3. HNeT and Neurocomputing

HNeT provides a major advance in the technology of machine pattern recognition. Operating to the complex domain, HNeT takes advantage of coherent addition to extend the capabilities of current neural technologies.

HNeT uses an input format consisting of one or more real-valued (or real-imaginary pairs) stimulus fields followed by one or more real-valued response fields. The response fields are required for training, but only provide the ground truthing for comparison and error calculation for the test/validation data.

HNeT is best suited for use with larger, complicated problems because HNeT relies on properties that converge only in the infinite limit. Thus, greater numbers of stimulus fields, memory elements and accurate training data will yield more stable results.

3.1 Coherence in Feature Space

In order for HNeT to operate effectively, the training set must be pre-processed as a whole. Either a feature space is pre-selected or is generated by filters/transforms in HNeT (Stimulus Conversion entries) from the stimulus fields. The distribution of this feature data will be renormalized to have zero mean and unit standard deviation.

The value of each feature needs to be mapped to a complex number, which has a polar angle, θ_n, so the range of values for each feature may need to be rescaled to the interval $(-\pi, \pi]$. Note that if operational data lies out of this range, then the training data has not provided a sufficiently representative set. When such an overlaps exist, ambiguities arise and the performance will suffer, perhaps significantly.

Coherent combination is applied over the entire training set, thereby reinforcing

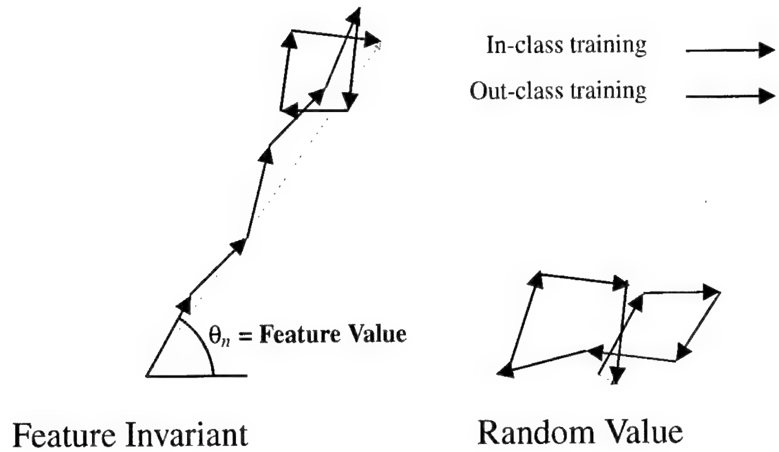

Figure 1: (Left) Common valued in-class features reinforce. (Right) Random valued features average to zero.

in-class invariant feature values, while random valued features average to zero, as illustrated in Figure 1. The length of the vector may be used to provide the confidence or weight of the image (not each feature), which may be controlled by varying the value of the response field. Multiple response fields may be used, operating in parallel, with each field defining a separate target class. For a given field k, each image is defined as out-class whenever the response $R(n,k) = 0$ and in-class otherwise. However, training is optimized for only one field, which can be selected by the user, so as to have an equal number of in-class and out-class vectors which allows the cancelation of invariants common to both.

Since vectors can vary in length, it is necessary to pair in-class and out-class vectors so as to approximately balance the contributions from each. It may happen that a feature may exhibit a random value for in-class vectors, but be invariant for the out-class. This may be acceptable for a closed problem, but for an open problem, where data not belonging to any of the pre-defined classes is allowed, this means that the out-class training set is not sufficiently representative.

How the coherent combination is implemented, when renormalization occurs, etc., is effectively controllable by the user. The number of epochs per training run, the learning rate, the amount of memory decay and mode are all parameters in HNeT (General settings), which may be tuned to the particulars of the given problem.

3.2 Feature Invariants

Once combined coherently, the vector lengths measure the reliability of the feature as an invariant for the class. A subset of the feature space template containing the best (longest) vectors are stored, one per memory element, providing the criteria to test new

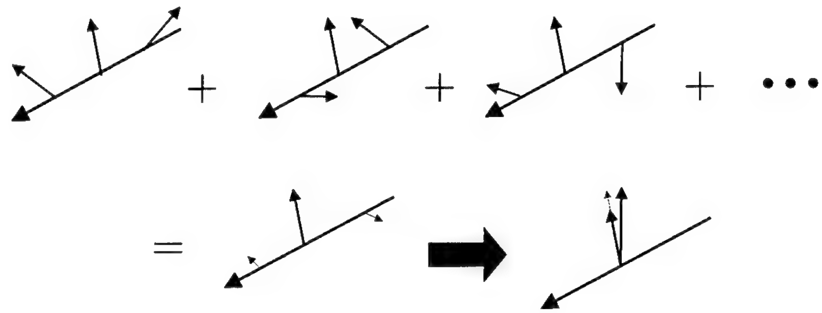

Figure 2: *Obtaining feature invariants through coherence (red). Best invariants form test criteria for new data (blue).*

data (Figure 2). The number of memory elements is user defined at the outset (Cortical Memory Elements settings). The amount of memory required is entirely dependent on the complexity of the classification problem being considered.

There will be some random-valued features that would average to zero in the infinite limit but, since only a finite number of samples are used for the training, exhibit a large degree of coherence. This is similar to speckle artifacts that appear in SAR imagery.

The effects of this feature space speckle can be reduced using HNeTs neural plasticity, whereby a user-defined percentage (Neural Plasticity settings) of the vectors in memory are examined for independence and pruned each pass in an attempt to improve the classification performance. A measure of classification is obtained by projecting the feature vectors of the test image onto the invariant feature vectors stored in memory.

Thus, since HNeT does not depend on pre-defined features, which may or may not separate the target classes for a given problem, but instead generates best features through coherence in the training data, this type of classifier is able to outperform most of the current pattern recognition products. Furthermore, since the methodology is an application offering the same kind of advantage as SAR does over *Real Aperture Radar* (RAR), many potential strengths and weaknesses of the technology should be understandable through analogy.

4. AND Contract Results

The AND Corporation was engaged by contract, SAR Automatic Target Recognition via a Holographic/Quantum Neural Net [22], to demonstrate the capability of their HNeT product for use as an ATR tool. The contract tasks included using unclassified 30 cm resolution X-band spotlight SAR images provided by DREO, i.e., the public MSTAR data set, for which pre-processing techniques would be developed to allow the data to be used by HNeT. Investigation on the effects of various parameters were to include minimum required resolution, target size, rotation, line-of-sight, speed of computation, signal to noise, contrast, partial obscuration and clutter.

Using the pixel magnitudes of the SAR imagery, it was determined that best results were obtained by logarithmically scaling the magnitudes and presenting the values in rectangular coordinates (rather than polar) to be transformed via a real-to-complex Fourier Transform. Centering on the target and rotating the images to a common target orientation were found to be critical to the processing.

Pixel averaging allowed the resolution to be reduced to 50 cm. This allowed the target areas to be extracted from the image in 64×64 chips so that a *Fast Fourier Transform* (FFT) could be implemented to reduce computation time. Furthermore, reducing the number of Fourier coefficients to 256, as well as eliminating the use of a windowing function (e.g., the Gabor transform) contributed to better processing speed.

Additionally, optimal HNeT performance was found to require approximately 1000 cortical memory elements with second order terms. An investigation of HNeT's neural plasticity over 100 optimization/regrowth cycles indicated 95% convergence of the parameters after 10 cycles. As well, higher performance was achieved using a binary predictor for each vehicle (each defining vehicle/not-vehicle classes) and selecting the vehicle class by vote.

Initially, the intention was to evaluate the ATR performance using two experiments applied to the MSTAR data [10]. The first was to separate the images at a single depression angle into training and testing sets that were separated by azimuthal angles of varying degrees. The second was to train on a single rotation cycle (azimuthal $0°$–$360°$) for each vehicle and test on the remainder.

The second experiment was replaced so as to take advantage of DARPA's division of the MSTAR data set, which identifies the $17°$ depression angle images as a training set and $15°$ as the corresponding test set. A third experiment was added, whereby the training set consisted of targets oriented within a $10°$ arc centered at $45°$, $90°$ and $180°$ orientation from the line-of-site. Testing sets cover $10°$ arcs of varying separation from the training arc.

In all cases, the performance criteria to be used was the Percent of Correct Classification, P_{cc}, thereby giving a single-valued result under a single operating condition. Since no confuser vehicles were included, all images were declared targets, i.e., no Other class existed, so $P_d = 1.0$ and the figure of merit is, in fact, the Percent of Correct Classification of Declared targets, $P_{cc|d}$. To clarify, the P_{cc} measure considers the rejection of a target image into the Other class as an error, while $P_{cc|d}$ removes all images declared as Other from the calculation.

4.1 Variable Azimuthal Angle

Although the experiment was first conducted using 10 vehicles at $15°$ depression angle, not all parameters had yet been optimized. Despite this, all results showed $P_{cc|d} \geq 88\%$ for up to a $30°$ azimuthal separation. When applied to the 4 vehicles available at $30°$

Table 1: Error Count for 30° depression angle with varying azimuthal separation.

Separation	T-72	ZSU-23-4	2S1	BTR-60	Errors	Images	$P_{cc\|d}$
1°	0	1	0	0	1	704	99.9 %
2°	0	0	0	0	0	766	100 %
3°	2	0	0	1	3	741	99.6 %
4°	0	0	0	0	0	775	100 %
5°	0	0	0	0	0	760	100 %
6°	0	0	0	0	0	752	100 %
7°	0	0	0	0	0	764	100 %
8°	1	0	0	0	1	750	99.9 %
9°	0	0	0	0	0	747	100 %
10°	0	0	0	0	0	772	100 %
11°	0	0	0	0	0	781	100 %
12°	0	0	0	0	0	794	100 %
13°	0	0	0	0	0	772	100 %
14°	0	0	0	0	0	756	100 %
15°	2	0	0	0	2	757	99.7 %
20°	0	0	0	0	0	789	100 %
25°	3	0	0	0	3	761	99.6 %
30°	0	1	0	0	1	709	99.6 %

and 45° depression angle, respectively, the results were in excess of 95%, as shown in Tables 1 and 2. However, no clear functional dependence on separation angle is evident.

Although these results provide high values of $P_{cc|d}$, the format of this experiment leaves many aspects ambiguous, such as the relative number of images used for each vehicle in the train and test sets. In other words, does the larger number of errors at low separation angles in the 45° results for the ZSU-23-4 indicate a problem with that class, or simply that more ZSU-23-4 images were used than the others? In fact, approximately twice as many ZSU-23-4 images are in the test set as compared to the other classes.

In order to compare the performance of classifiers, it is important to choose an experiment for which $P_{cc|d}$ varies significantly over the domain of the varying parameters, and somewhat smoothly with respect to the step size of variation. Clearly, this is not the case for this experiment.

4.2 Variable Depression Angle

The MSTAR data is disseminated with the 17° depression angle images identified as training files and the 15° depression angle as a test set, which is consistent with the recommendations of the MSTAR evaluation [1]. Based on this, a second experiment was designed to train 10 binary classifiers on the 17° data and test at 15°. All available data in each set was used and the images pre-processed according to four methods: the

Table 2: *Error Count for 45° depression angle with varying azimuthal separation.*

Separation	T-72	ZSU-23-4	2S1	BTR-60	Errors	Images	$P_{cc\|d}$
1°	0	1	0	0	1	709	99.9 %
2°	1	3	5	4	13	796	98.4 %
3°	1	10	3	5	19	760	97.5 %
4°	3	9	2	5	19	791	97.6 %
5°	1	10	3	6	20	773	97.4 %
6°	2	4	3	6	15	792	98.1 %
7°	1	4	5	3	13	769	98.3 %
8°	5	8	6	4	23	785	97.1 %
9°	4	7	2	5	18	770	97.7 %
10°	5	10	5	4	24	797	97.0 %
11°	2	6	4	3	15	751	98.0 %
12°	3	3	5	4	15	772	98.1 %
13°	6	2	8	6	22	775	97.2 %
14°	4	2	14	3	23	778	97.0 %
15°	6	11	6	5	28	797	96.5 %
20°	2	9	5	8	24	797	97.0 %
25°	5	10	12	14	41	779	94.7 %
30°	5	3	11	15	34	772	95.6 %

raw pixel magnitudes of the 64×64 target segments, logarithmic scaling of the pixel magnitudes, rotation of the chips to a common alignment and a combined alignment plus logarithmic scaling. Orientation of the vehicle heading is extracted from the ground truthing for images in both the training and testing sets. The error count of false negatives for each binary classifier is reported in Table 3.

Although a degradation from the raw image results occurs when logarithmic scaling to the BMP-2 and BTR-70 images is applied, an improvement is obtained when the logarithmic scaling is combined with alignment. Because the entire imagery was used in a single implementation of each preprocessing method, rather than developing a confidence interval over several runs on subsets, it is difficult to identify the reason for the anomaly. Again, although the overall trend is more apparent here, the range of results does not vary sufficiently to allow for good comparison of classifier ability.

4.3 Line-of-Sight Orientation

Although not a standard evaluation experiment, the third experiment does provide sufficient variation for a more meaningful analysis of the classifier performance, even though the orientation variable is usually addressed during preprocessing before classification occurs. For training, all images at both 15° and 17° and within a 10° arc, centered at orientations of 45°, 90° or 180° from the radar *Line-of-Sight* (LOS), are selected. Test images are selected based on azimuthal separation from the training set,

Table 3: *Error count for HNeT classifier for MSTAR images at 15° depression angle, under several preprocessing methods. Training data used 17° depression angle.*

	Raw	Log	Aligned	Aligned + Log
T-72	11	1	6	0
BTR-70	13	21	9	1
BMP-2	20	43	13	6
ZSU-23-4	2	0	5	0
Zil-131	14	3	10	0
T-62	5	1	3	0
D-7	1	0	1	0
BRDM-2	2	1	2	0
2S1	45	18	11	0
BTR-60	6	3	4	2
Total Errors	119	91	64	9
Total Records	5392	5392	5392	5392
$P_{cc\|d}$	97.79 %	98.31 %	98.81 %	99.83 %

as shown in Table 4.[1]

As before, however, the reporting method does not provide any insight into the actual number of images used for each class. In fact, both training and testing sets are heavily weighted with T-72 images, which skews the results with a bias toward the T-72 solution. Typically, $1/2$ the images are T-72, $1/4$ are BMP-2 and the remaining $1/4$ are from the 8 other classes. Thus, the high $P_{cc|d}$ at large separation angles may appear to indicate successful classification when, in fact, only the T-72 classifier is performing well. In particular, the BMP-2 results which show more than 18 errors indicate that this classifier is performing worse than a *random* classifier. Clearly, the training/testing design is in need of modification for this experiment.

5. Standard Evaluation of HNeT

Analysis of the results produced from the contract have shown several limitations to the information thereby generated. As such, it was determined that the performance of the HNeT system needed to be evaluated according to standard techniques found in the literature [23, 7, 24]. Furthermore, a careful examination of the design of the preprocessing techniques needed to be carried out.

[1]The AND Corporation report gives 5 errors for the BMP-2 at LOS $+ 45°$ and 45° separation, but the full data indicates 37 errors, as shown here.

Table 4: *Error count for HNeT binary classifiers of each MSTAR class according to azimuthal separation from the training set. Training set taken over a 10° arc, centered at 45°, 90° and 180° orientation from the radar LOS.*

LOS $+45°$	T-72	BTR-70	BMP-2	ZSU-23-4	Zil-131	T-62	D-7	BRDM-2	2S1	BTR-60	Errors	Images	$P_{cc\|d}$
5°	0	0	0	0	0	0	0	0	0	0	0	275	100 %
10°	0	2	0	0	0	0	0	0	1	0	3	270	98.9 %
15°	0	0	2	0	0	0	0	0	3	0	5	255	98.0 %
20°	0	1	2	0	1	0	0	0	1	0	5	210	97.7 %
25°	1	1	6	0	0	0	0	0	1	0	9	279	96.8 %
30°	0	1	19	0	0	1	0	0	4	1	26	278	90.6 %
35°	2	0	21	2	1	1	0	0	5	2	34	248	86.3 %
40°	0	0	28	1	0	1	0	0	2	2	34	173	80.3 %
45°	0	2	37	3	5	4	3	0	9	5	68	221	69.2 %

LOS $+90°$	T-72	BTR-70	BMP-2	ZSU-23-4	Zil-131	T-62	D-7	BRDM-2	2S1	BTR-60	Errors	Images	$P_{cc\|d}$
5°	0	0	0	0	0	0	0	0	0	0	0	242	100 %
10°	0	0	1	0	0	0	0	0	0	0	1	273	99.6 %
15°	1	0	5	0	0	2	1	0	2	0	11	266	95.9 %
20°	0	0	9	4	0	1	1	0	1	0	16	236	93.2 %
25°	0	0	11	2	4	0	0	0	2	0	19	231	91.8 %
30°	1	2	18	3	2	8	2	0	4	0	40	263	84.8 %
35°	1	4	20	3	2	5	4	0	7	0	46	276	83.3 %
40°	5	3	22	4	2	2	3	0	4	4	49	263	81.4 %
45°	3	2	31	6	3	3	2	0	5	1	56	277	79.8 %

LOS $+180°$	T-72	BTR-70	BMP-2	ZSU-23-4	Zil-131	T-62	D-7	BRDM-2	2S1	BTR-60	Errors	Images	$P_{cc\|d}$
5°	0	0	0	0	0	0	1	0	0	0	1	252	99.6 %
10°	0	3	0	0	0	0	0	0	0	0	3	266	98.9 %
15°	0	1	1	0	1	0	0	1	0	0	4	231	98.3 %
20°	0	1	7	1	0	0	0	0	1	0	10	223	95.5 %
25°	1	2	9	0	3	3	0	0	1	0	19	240	92.1 %
30°	0	3	16	0	3	5	0	1	4	0	32	281	88.6 %
35°	0	6	13	0	2	3	0	1	5	0	30	278	89.2 %
40°	0	2	19	0	7	5	0	1	4	0	38	277	86.3 %
45°	0	7	15	2	6	7	0	1	6	0	44	274	83.9 %

Table 5: Confusion matrix of MSTAR 15° imagery for HNeT 10-vehicle classifier trained on 17° data. Correctly classified images contribute to the diagonal elements while misclassifications appear off-diagonal.

	2S1	BMP-2	BRDM-2	BTR-60	BTR-70	D7	T-62	T-72	Zil-131	ZSU-23-4	
2S1	**274**	0	0	0	0	0	0	0	0	0	100.00%
BMP-2	0	**581**	0	0	0	0	0	5	1	0	98.98%
BRDM-2	0	0	**274**	0	0	0	0	0	0	0	100.00%
BTR-60	0	1	0	**193**	0	0	0	1	0	0	98.97%
BTR-70	1	0	0	0	**195**	0	0	0	0	0	99.49%
D7	0	0	0	0	0	**274**	0	0	0	0	100.00%
T-62	0	0	0	0	0	0	**273**	0	0	0	100.00%
T-72	0	0	0	0	0	0	0	**2771**	0	0	100.00%
Zil-131	0	0	0	0	0	0	0	0	**274**	0	100.00%
ZSU-23-4	0	0	0	0	0	0	0	0	0	**274**	100.00%
											99.83%

5.1 Meeting Standards

First, a complete separation of the training and testing phases was necessary. During this investigation, it was determined that the HNeT neural plasticity, i.e., the optimization of the selection of invariants, had been made relative to the performance of classification for the test set, meaning that the classifier predictions were being influenced by knowledge of the actual target class of the test imagery. The HNeT settings were, therefore, adjusted to run the optimization against the training set and the test data was not even selected until after training was complete. Despite this, there were no significant changes to the HNeT results for the AND experiments when those trials were redone with these now independent settings.

Validation with data independent of the training process and fully sanitized of ground truth resulted in near-identical results. Because the neural plasticity optimization is not a deterministic process, identical results could not be obtained and should not be expected. The non-repeatability of this process does, however, emphasize the need to produce confidence intervals rather than single-point results, and that training trials need to be repeated as well as those for testing.

Returning to the AND data used to generate Table 3, the information can be presented in a more meaningful way by used of the standard confusion matrix [24]. Each row corresponds to the set of imagery for the indicated class. Cell entries give the number of images from the row set that have been classified as the vehicle type indicated by the column header.

Thus, the matrix form provides a greater indicator of what mistakes the classifier is making as well as the relative number of images in each class. In addition, the results are summarized by effectiveness on each class of imagery, with an overall metric, $P_{cc|d}$ indicating the classifier success rate.

Table 5 gives the confusion matrix of MSTAR imagery for the HNeT classifier with

Table 6: *Confusion matrix for CIS MMSE classifier [25] on MSTAR 15° imagery, with training at 17°.*

	2S1	BMP-2	BRDM-2	BTR-60	BTR-70	D7	T-62	T-72	Zil-131	ZSU-23-4	
2S1	**262**	0	0	0	0	0	4	8	0	0	95.62%
BMP-2	0	**581**	0	0	0	0	0	6	0	0	98.98%
BRDM-2	5	3	**227**	1	0	14	3	5	4	1	85.31%
BTR-60	1	0	0	**193**	0	0	0	0	0	1	98.97%
BTR-70	4	5	0	0	**184**	0	0	3	0	0	93.88%
D7	2	0	0	0	0	**271**	1	0	0	0	98.91%
T-62	1	0	0	0	0	0	**259**	11	2	0	94.87%
T-72	0	0	0	0	0	0	0	**582**	0	0	100.00%
Zil-131	0	0	0	0	0	0	2	0	**272**	0	99.27%
ZSU-23-4	0	0	0	0	0	0	0	1	0	**271**	98.91%
											97.18%

preprocessing by logarithmic scaling of pixel values and common target orientation. Here, it can easily be seen that the BMP-2 error is again biased to the T-72, a problem originally traceable only in the Line-of-Sight experiment when using the AND method of reporting.

Table 6 provides a similar confusion matrix published by O'Sullivan *et al.* [25] from the *Center for Imaging Science* (CIS) using a *minimum mean-squared error* (MMSE) estimator, a single-rule based classifier, against conditionally Gaussian signal models of the vehicle classes. It is clear this method also yields a significant bias towards T-72 classification, but other secondary errors are being made that HNeT avoids.

5.2 Evaluating Classifier Performance

Standard methods for evaluating the performance of a classifier include the use of Confusion Matrices and the generation of Receiver Operating Characteristic curves [26, 7, 25]. Specifically, the standard evaluation suggested by DARPA/WL applied to the MSTAR data set [1] provides a baseline with which to compare. Using three vehicle classes for which multiple vehicles are present in the MSTAR data set, the classifier is trained on imagery at 17° depression angle for a single vehicle from each class and then tested on imagery at 15° depression for all the vehicles.

DARPA/WL used 3 BMP-2 (serial numbers c21, 9563 and 9566) infantry fighting vehicles, 4 BTR-70 (s/n c72, c70, c73 and c71) armored personnel carriers and 3 T-72 (s/n 132, 812 and s7) main battle tanks (see Figure 3). All the BMP-2 and T-72 images were included in the released set but, unfortunately, the public release data included only one BTR-70, s/n c71. This means that although the ROC curves for BTR-70s cannot be directly compared, the other classes can. Furthermore, the training BTR-70 used by DARPA/WL is not the one included in the public release. Despite this, recalculating the DARPA/WL results without the extra BTR-70s still does allow for the comparison of Confusion Matrices, so long as the potential differences due to the surrogate vehicle are recognized. In particular, the removal of the 3 extra BTR-70

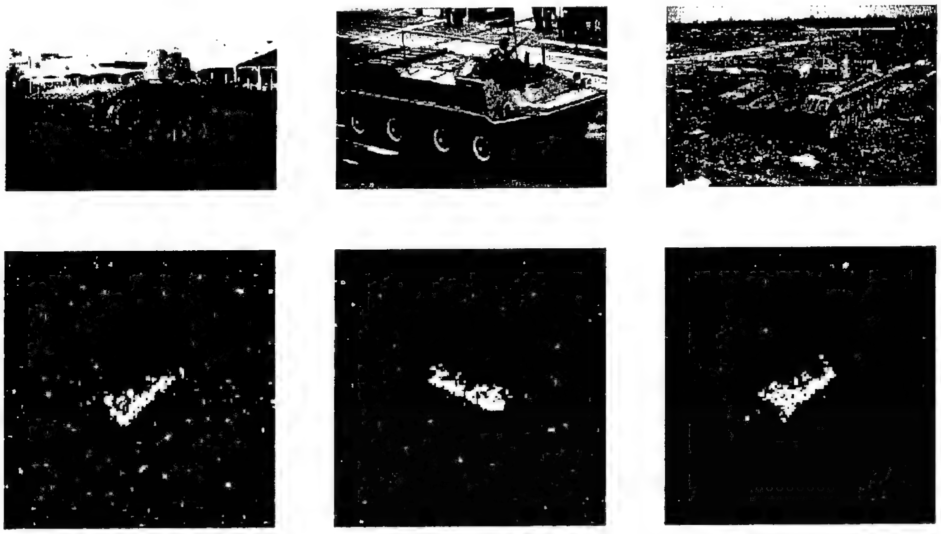

Figure 3: Photo (top) and SAR (bottom) images of vehicles used for MSTAR evaluation criteria: (left) BMP-2, (centre) BTR-70 and (right) T-72.

vehicles shifts the operating point of the confusion matrix from $P_d = 0.90$ to $P_d = 0.923$.

5.3 HNeT vs. MSTAR Classifier Comparison

Subject to the use of BTR-70 s/n c71 instead of c72 as the imagery source for that vehicle class, the HNeT classifier was trained as indicated by the MSTAR standard. At 17° depression angle, the training set was composed of 233 images of BMP-2 s/n c21, 233 images of BTR-70 s/n c71 and 232 images of T-72 s/n 312, along with 298 images of "slicy," a non-vehicle target (Figure 4) included so as to improve the definition of the out-class. As with the previous HNeT experiments, the image chips were aligned according to the ground truth heading and the magnitudes of the central 64×64 pixel block extracted. The resulting data set was used to train three binary classifiers, one corresponding to each vehicle class.

Figure 4: Photo (left) and SAR (right) images of the "slicy" non-vehicle target.

Table 7: Confusion matrices for standard MSTAR evaluation with training vehicles indicated by asterisks. (Left) MSTAR baseline classifier [1] modified for public release data. (Right) HNeT classifier.

MSTAR baseline Confusion Matrix ($P_d = 0.92$)

	BMP-2	BTR-70	T-72	Other
BMP-2 (9563)	**161**	18	0	*16*
BMP-2 (9566)	**150**	31	0	*15*
BMP-2 (c21)*	**175**	8	0	*13*
BTR-70 (c72)*	0	**270**	0	*3*
T-72 (132)*	1	2	**188**	*5*
T-72 (812)	13	35	**112**	*35*
T-72 (s7)	8	28	**131**	*24*
				$P_{cc\|d}$= **0.8918**

HNeT Confusion Matrix ($P_d = 0.92$)

	BMP-2	BTR-70	T-72	Other
BMP-2 (9563)	**185**	0	4	*6*
BMP-2 (9566)	**167**	2	12	*15*
BMP-2 (c21)*	**191**	1	0	*4*
BTR-70 (c71)*	0	**195**	0	*1*
T-72 (132)*	0	0	**194**	*2*
T-72 (812)	30	2	**122**	*41*
T-72 (s7)	9	8	**138**	*36*
				$P_{cc\|d}$= **0.9460**

HNeT was configured to perform a 2D FFT on the stimulus fields (as a 64×64 array) and with 2000 memory elements. Training occurred for 50 cycles for 4 epochs with a neural plasticity at 50% optimization and based on minimizing the mean absolute error between HNeT's recall and the desired response.

Once the classifiers were trained, the 15° depression angle imagery was pre-processed in a manner identical to what was done for the training data, this time including all vehicle imagery within each class. This test data was input to the HNeT classifiers as a validation set and the response recall calculated and exported. These results were compared, using a *winner takes all* (WTA) strategy, i.e., each image is classified according to the classifier with the highest HNeT response value.

At this point, a threshold value was introduced, whereby any image having all three responses below the threshold was rejected and moved to the "Other" class. The value of the threshold was adjusted to 0.5275 in order to yield a rejection ratio of 0.077, corresponding to $P_d = 0.923$, the operating point for the modified MSTAR baseline.

Table 7 compares the MSTAR baseline and HNeT Confusion Matrices for the 3-class multi-vehicle evaluation. Correct classification is indicated in bold, resulting in a $P_{cc|d}$ of 89% for the MSTAR baseline and 95% for HNeT. Note that for both classifiers, all BTR-70 images are either correctly classified or designated as Other and since the MSTAR baseline has more images, the $P_{cc|d}$ comparison will be biased in MSTAR's favour. It is also apparent that the two types of classifiers are making very different errors when they happen. The MSTAR method has both BMP-2 and T-72 images being misclassified as BTR-70s, while the HNeT results show them being confused with each other, and not the BTR-70. While the different vehicle used may have some contribution, the major effect is expected to be the difference in features used.

To produce ROC curves, the MSTAR experiment used two confuser vehicles to generate false alarms, an M-109 self-propelled howitzer and an M-110 self-propelled howitzer, neither of which are part of the public release. Instead, for HNeT the imagery of the six remaining public release vehicles was used, namely the 2S1 self-propelled

howitzer, the D7 bulldozer, the T-62 main battle tank, the BTR-60 armored personnel carrier, the Zil-131 truck and the ZSU-23-4 self propelled anti-aircraft gun.

Since the confusers for the MSTAR baseline are both of similar design, i.e. a large gun on a tracked chassis, it can be expected that most of the confusion will be biased to one class, with the others less affected. The confusers used for the HNeT experiment offer vehicles similar to each of the three classes. Furthermore, the MSTAR confusers are American design, while all the targets are former USSR equipment. The HNeT confusers are, like the targets, all of Soviet origin.

These factors suggest that the same classifier should render more false alarms for the HNeT experiment than the MSTAR. Nevertheless, as seen in Figure 5, the HNeT classifier outperforms the MSTAR baseline since all three HNeT curves are closer to the ideal operating point at $(0,1)$.

For each binary HNeT classifier, the ROC curve was generated by setting the threshold at a value larger than the maximum HNeT response for all test images, both targets and confusers. Here, all images fall below the threshold and are classified as Other, thus $P_d = 0$ and $P_{fa} = 0$. As the threshold is stepped downward, a parametric curve is produced as target images above the threshold increase P_d and high values for confuser images increase P_{fa}. Once the threshold is below the minimum HNeT response, all images are classified, i.e., $(P_d, P_{fa}) = (1,1)$.

If there had existed one or more values of the threshold where all the target images fall above the threshold and all the confusers below, then $(P_d, P_{fa}) = (1,0)$ and the classifier would have had an ideal operating point *for this set of images*. Since this is not the case, proximity to $(1,0)$ is a measure of performance, as is the total area under the ROC curve.

The diagonal line in Figure 5 represents a random classifier, where the output class is completely independent of the type of image given as input. Thus, the chances for an image being classified as the target class are the same whether that image is in-class or out-class, i.e. $P_d = P_{fa}$. The actual probability of that classification, P_{rdm}, parameterizes the line the same way as the threshold does the ROC curves, so we have the straight line, $(P_d, P_{fa}) = (P_{\mathrm{rdm}}, P_{\mathrm{rdm}})$.

In addition to the MSTAR baseline, the author became aware during the writing of this report that in a more recent publication from the *Computational NeuroEngineering Laboratory* (CNEL) at the University of Florida, Zhao *et al.* evaluated their SVM classifier using the MSTAR standard and were also limited to using only the public release imagery. Essentially, the same decisions, arrived at independently, of how to adapt the evaluation method were made as described herein, except that CNEL made no effort to generate a confusion matrix at $P_d = 0.92$.

Figure 5 compares the ROC curves generated by the MSTAR baseline, CNEL's SVM and the HNeT classifiers. Again, the HNeT and CNEL curves for the BTR-70 must be

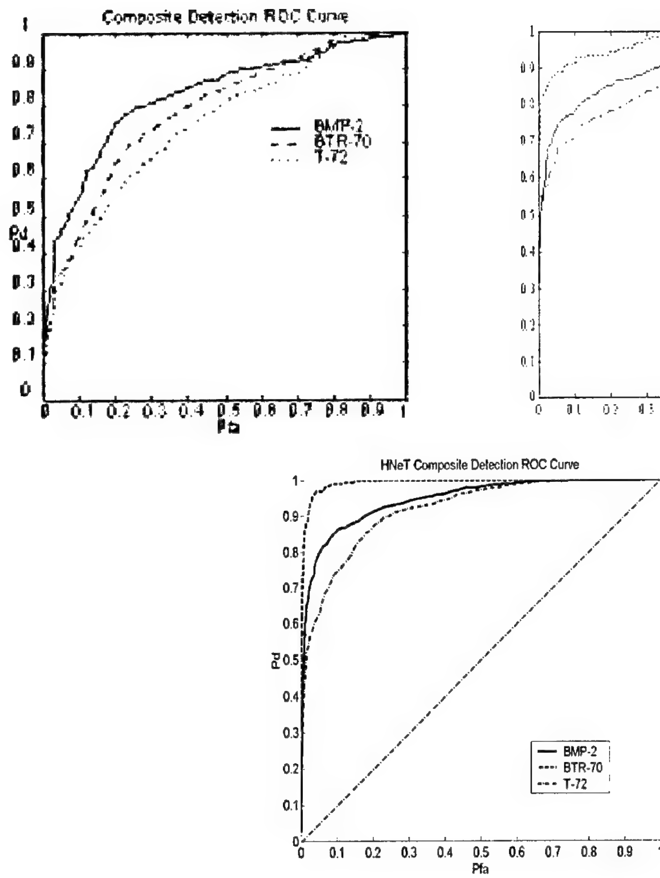

Figure 5: *ROC curves for the MSTAR standard evaluation of (top left) the MSTAR baseline classifier [1], (top right) the CNEL SVM classifier [21], and (bottom) the HNeT classifier. The HNeT and CNEL BTR-70 classifiers are tested on the same vehicle as training. The diagonal corresponds to a random classifier.*

interpreted differently compared to either the other results because these classifiers have not been applied to test vehicles other than the ones they were trained upon and so the curves are biased with a higher P_d expectation than the others. If applied to the withheld imagery, it should be expected that the BTR-70 curves in question would be shifted to the same region as their corresponding BMP-2 and T-72 curves. In any case, with either measure, the HNeT results clearly show a better performance than either the MSTAR baseline or CNEL SVM classifier.

The CNEL publication also generated $P_{fa} - P_{cc|d}$ ROC curves for each of the vehicle variants, whereby the classifier success is measured not only according to the rate of declaring target images, but of both declaring and correctly classifying them. In such plots, a random classifier follows a straight line from $(0,0)$ to $(1, 1/N)$ where N is the number of target classes. Figure 6 compares the CNEL SVM and HNeT classifiers using these ROC curves. Again, a significant improvement in performance is seen for HNeT.

6. Conclusion

Significant insight has been gained into the methodology behind AND Corporation's HNeT software. The company's claims and demonstrations of superior performance compared to most current classifier technology are now understood and verified. Sufficient understanding of their method of extracting feature invariants via coherence, especially through analogy with the capabilities of SAR over RAR, generates confidence in proceeding to use HNeT as an engine for an ATR system. Furthermore, this understanding has identified limitations in HNeT which could be overcome with future research efforts.

The experiments performed by AND Corporation under contract with DRDC demonstrate both the success and limitations of their classifier, as well as determining optimal configuration of HNeT for these experiments. Unfortunately, the design of the experiments is not ideal, allowing bias to skew some results, such as the influence of the number of available images.

By using the standard evaluation methods published by DARPA/WL for the MSTAR data collection, a more reliable measure of the HNeT classifier performance has been obtained. Confusion matrices and ROC curves tailored for the publicly released MSTAR data collection have been generated and compared to an MMSE classifier from CIS, the MSTAR baseline results and an SVM classifier from CNEL. These all show a significant increase in performance by using the HNeT classifier over the others, specifically with the ROC curves for HNeT reducing the region of error from MSTAR by about half, as evidenced by the $P_d = 0.92$ Confusion Matrix yielding an improvement from $P_{cc|d} = 0.89$ for the MSTAR baseline to $P_{cc|d} = 0.95$ for HNeT.

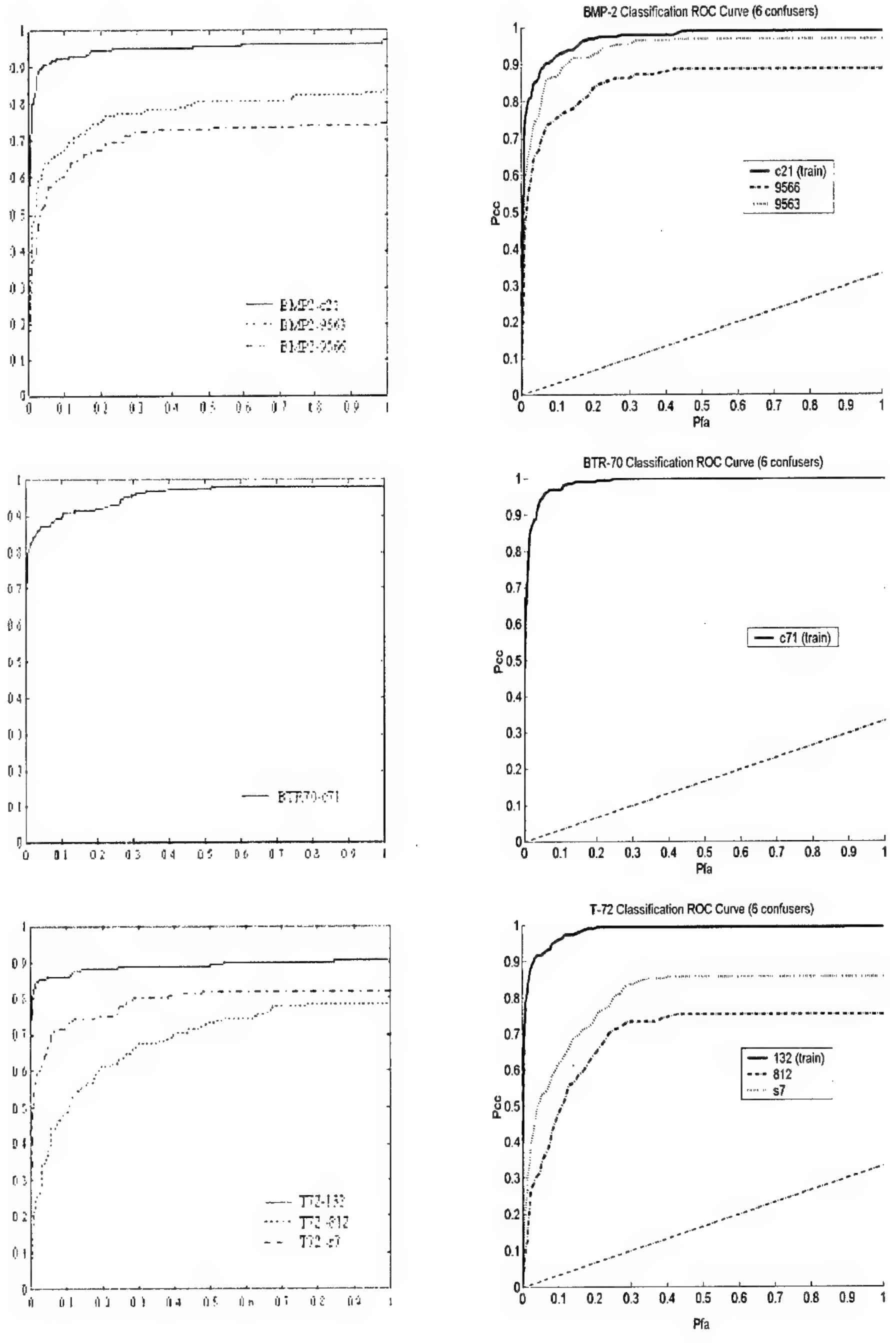

Figure 6: *$P_{cc|d}$ versus P_{fa} ROC curve break-out for (top) BMP-2 variants, (middle) BTR-70 and (bottom) T-72 variants, comparing (left) the CNEL SVM classifier [21] and (right) the HNeT classifier.*

References

1. T.D. Ross, S.W. Worrell, V. Velten, J.C. Mossing and M.L. Bryant (1998). Standard SAR ATR Evaluation Experiments Using the MSTAR Public Release Data Set. In *Proc. SPIE, Vol. 3370, Algorithms for SAR Imagery V*, p. 566. Orlando, USA.

2. R.O. Duda and P.E. Hart (1973). Pattern classification and scene analysis, New York, USA: Wiley.

3. D.S. Sivia (1996). Data Analysis: a Bayesian Tutorial, New York, USA: Oxford University Press.

4. R. Haralick (1986). Computer Vision Theory: The Lack Thereof. *Computer Vision, Graphics, and Image Processing*, **36**, 372.

5. R. Jain and T. Binford (1991). Ignorance, Myopia, and Naivete in Computer Vision Systems. *CVGIP: Image Understanding*, **53**(1), 112.

6. T. Pavlidis (1992). Why Progress in Machine Vision Is So Slow. *Pattern Recognition Lett.*, **13**, 221.

7. D.E. Dudgeon (2000). ATR Performance Modeling and Estimation. *Digital Signal Processing*, **10**(4), 269.

8. M. Boillot and J. Principe (1998). An Algorithm Development and Support Database for MSTAR. In *Proc. of the Image Understanding Workshop, Vol. 2*, p. 821. Pacific Grove, USA.

9. ALPHATECH (1998). MSTAR PET. (Online) ALPHATECH, Inc. `http://www.alphatech.com/secondary/techpro/projects/mstar/PET.html` (August 1, 2001).

10. J. Sutherland (1999). SAR Automatic Target Recognition via a Holographic/Quantum Neural Net. (Contract W7714-8-0200/001/SV). Toronto: AND Corporation.

11. K.L. Sala (1997). Image Classification by Neural Networks Using Moment Invariant Feature Vectors. (CRC Report No. 97-002). Ottawa: Communications Research Centre, Industry Canada.

12. D.G. Arnold and K. Sturtz (2000). Geometric Object/Image Relations for RADAR. In *Proc. SPIE, Vol. 4053, Algorithms for SAR Imagery VII*, p. 538. Orlando, USA.

13. E. Ertin and L.C. Potter (2000). Polarimetric Classification of Scattering Centers Using *M*-ary Bayesian Decision Rules. *IEEE Trans. Aerospace and Electronic Systems*, **36**(3), 738.

14. Y. Tessier (2000). Evaluation of Decision Rule Based Neural Net on SpotSAR Imagery of Military Ships. (990001532). Montréal: Lockheed Martin Canada.

15. T.-W. Lee (1998). Independent Component Analysis: Theory and Applications, Boston, USA: Kluwer Academic Publishers.

16. A. Hyvärinen and E. Oja (2000). Independent Component Analysis: Algorithms and Applications. *Neural Networks*, **13**(4–5), 411.

17. H.-C. Chiang, R.L. Moses and L.C. Potter (2000). Model-Based Classification of Radar Images. *IEEE Trans. on Information Theory*, **46**(5), 1842.

18. J. Hwang and S. Kung and M. Niranjan and J. Principe (1997). The Past, Present and Future of Neural Networks for Signal Processing. *IEEE Signal Processing Magazine*, **14**(6), 28.

19. J. Principe and M. Kim and J. Fisher (1998). Target Discrimination in Synthetic Aperture Radar (SAR) Using Artificial Neural Networks. *IEEE Trans. on Image Processing*, **7**(8), 1136.

20. T.H. Kerr (1998). Critique of Some Neural Network Architectures and Claims for Control and Estimation. *IEEE Trans. Aerospace and Electronic Systems*, **34**(2), 406.

21. Q. Zhao, J. Principe, V. Brennan, D. Xu and Z. Wang (2000). Synthetic Aperture Radar Automatic Target Recognition with Three Strategies of Learning and Representation. *Optical Engr.*, **39**(5), 1230.

22. Public Works and Government Services Canada (1999). Contract for *SAR Automatic Target Recognition via a Holographic/Quantum Neural Net.* PWGSC 007SV.W7714-8-0200. Public Works and Government Services Canada.

23. J.C. Mossing and T.D. Ross (1998). An Evaluation of SAR ATR Algorithm Performance Sensitivity to MSTAR Extended Operating Conditions. In *Proc. SPIE, Vol. 3370, Algorithms for SAR Imagery V*, p. 554. Orlando, USA.

24. S.G. Alsing (2000). The Evaluation of Competing Classifiers. Ph.D. thesis. Air Force Institute of Technology. Wright-Patterson AFB, OH.

25. J.A. O'Sullivan, M.D. DeVore, V. Kedia and M.I. Miller (2001). SAR ATR Performance Using a Conditionally Gaussian Model. *IEEE Trans. Aerospace and Electronic Systems*, **37**(1), 91.

26. T.D. Ross and J. Mossing (1999). MSTAR Evaluation Methodology. In *Proc. SPIE, Vol. 3721, Algorithms for SAR Imagery VI*, p. 705. Orlando, USA.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.) Defence Research Establishment Ottawa 3701 Carling Avenue Ottawa, ON K1A 0Z4	2. SECURITY CLASSIFICATION (overall security classification of the document, including special warning terms if applicable) UNCLASSIFIED

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C or U) in parentheses after the title.)

Automatic Target Recognition Using HNeT: an investigation of Holographic/Quantum Neural Technology (U)

4. AUTHORS (Last name, first name, middle initial)

English, Ryan A.

5. DATE OF PUBLICATION (month and year of publication of document) December 2001	6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.) 23	6b. NO. OF REFS (total cited in document) 26

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.)

Defence Research Establishment Ottawa
3701 Carling Avenue
Ottawa, ON K1A 0Z4

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant) 3db31	9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written)
10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.) DREO TM 2001-080	10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned this document either by the originator or by the sponsor)

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
(   ) Distribution limited to defence departments and defence contractors; further distribution only as approved
(   ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
(   ) Distribution limited to government departments and agencies; further distribution only as approved
(   ) Distribution limited to defence departments; further distribution only as approved
(   ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)

Unlimited

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

(U) With the release of the Moving and Stationary Target Acquisition and Recognition (MSTAR) public data set, high quality Synthetic Aperture Radar (SAR) imagery of military ground vehicles has been made accessible to the entire research community. Furthermore, standard methods for evaluating classifier results on this data set have been created and released. Using these tools, we reconsider a previously contracted application of AND Corporation's Holographic/Quantum Neural Technology (HNeT) classifier, performing brief analyses of the way HNeT selects features for Automatic Target Recognition (ATR) purposes, the methodology used in the contract and their results, as well as obtaining new results that comply with the MSTAR standard evaluation criteria. These results provide measures of performance for the HNeT classifier using Confusion Matrices and Receiver Operating Characteristic (ROC) curves, that are used to compare with the open literature performance of the MSTAR baseline and two other classifiers, from which we conclude that HNeT outperforms the other three and provides improved ATR.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Automatic Target Recognition, ATR, feature invariants, HNeT, classifier, MSTAR, SpotSAR imagery

Defence R&D Canada

is the national authority for providing
Science and Technology (S&T) leadership
in the advancement and maintenance
of Canada's defence capabilities.

R et D pour la défense Canada

est responsable, au niveau national, pour
les sciences et la technologie (S et T)
au service de l'avancement et du maintien des
capacités de défense du Canada.

DEFENCE R&D DÉFENSE

www.drdc-rddc.dnd.ca